arXiv:1512.06061v1 [cs.LG] 18 Dec 2015

Asymptotic Behavior of Mean Partitions in Consensus Clustering

Brijnesh J. Jain
Technische Universität Berlin, Germany
e-mail: brijnesh.jain@gmail.com

Although consistency is a minimum requirement of any estimator, little is known about consistency of the mean partition approach in consensus clustering. This contribution studies the asymptotic behavior of mean partitions. We show that under normal assumptions, the mean partition approach is consistent and asymptotic normal. To derive both results, we represent partitions as points of some geometric space, called orbit space. Then we draw on results from the theory of Fréchet means and stochastic programming. The asymptotic properties hold for continuous extensions of standard cluster criteria (indices). The results justify consensus clustering using finite but sufficiently large sample sizes. Furthermore, the orbit space framework provides a mathematical foundation for studying further statistical, geometrical, and analytical properties of sets of partitions.

Contents

1. Introduction **3**

2. Geometry of Partition Spaces **4**
2.1. Partitions 4
2.2. The Orbit Space of Partitions 4
2.3. Metric Structures 5

3. Fréchet Consensus Clustering **6**
3.1. Fréchet Functions 6
3.2. Consistency by Metric Structures 7
3.3. Consistency by Stochastic Programming 8
3.4. A Central Limit Theorem 9

4. Generalized Criteria for Comparing Partitions **9**
4.1. Criteria based on Counting Pairs 10
4.2. Criteria based on Cluster Matchings 12
4.3. Criteria based on Information Theory 13

5. Conclusion **14**

A. Bhattacharya-Patrangenaru Consistency **14**

B. Proofs **15**
B.1. Proof of Theorem 2.1 15
B.2. Proof of Theorem 3.3 16
B.3. Proof of Theorem 3.4 19
B.4. Proofs of Statements in Section 4 19

1. Introduction

Clustering is a standard technique for exploratory data analysis that finds applications across different disciplines such as computer science, biology, marketing, and social science. The goal of clustering is to group a set of unlabeled data points into several clusters based on some notion of dissimilarity. Inspired by the success of classifier ensembles, consensus clustering has emerged as a research topic [9, 22]. Consensus clustering first generates several partitions of the same dataset. Then it combines the sample partitions to a single consensus partition. The assumption is that a consensus partition better fits to the hidden structure in the data than individual partitions.

One standard approach of consensus clustering combines the sample partitions to a mean partition [4, 5, 6, 7, 10, 14, 18, 20, 21]. A mean partition best summarizes the sample partitions with respect to some (dis)similarity function. A natural question is the choice of sample size. If the sequence of mean partitions fails to converge stochastically for growing sample size n, then picking a reasonable value for n becomes an additional parameter selection problem. Otherwise, if the mean partitions converge stochastically to an expected partition, the problem of selecting a sample size n simplifies to the problem of selecting a sufficiently large n, because we have high confidence that nothing unexpected will happen when sampling further partitions. In other words, stochastic convergence justifies the common practice to draw finite but sufficiently large sample sizes.

Though there is an extensive literature on consensus clustering [22], little is known about the asymptotic behavior of the mean partition approach. Topchy et al. [19] studied the asymptotic behavior of the mean partition approach under the following simplifying assumptions:

(A1) The underlying distance is a semi-metric.[1]

(A2) Partitions are hard (crisp) partitions.

(A3) The expected partition is unique.

(A4) Partitions strongly concentrate on the expected partition.

In this contribution, we study the asymptotic behavior of the mean partition approach without drawing on assumptions (A2)–(A4). We show (i) consistency of the mean partitions, (ii) strong consistency of the variations, and (iii) a modified version of the Central Limit Theorem for mean partitions. We present two variants of results (i) and (ii). The first variant assumes that partitions form a compact metric space. The second variant requires the Euclidean space as ambient space and assumes that partitions are compared by a continuous cluster criterium. We also draw on continuity of the cluster criterium for showing result (iii). Since standard criteria for comparing partitions are defined on the discrete space of hard partitions, we present examples of their continuous extensions. We can apply the generalized standard criteria to soft partitions and we can analyze the asymptotic behavior of the mean partition approach for hard and soft partitions in a unified manner.

[1] A semi-metric satisfies all axioms of a metric, but not necessarily the triangle inequality.

The basic idea to derive the results is to represent partitions as points of a geometric space, called orbit space. Orbit spaces are well explored, possess a rich mathematical structure and have a natural connection to Euclidean spaces [3, 11, 16]. For the first variant of results (i) and (ii), we link the consensus function of the mean partition approach to Fréchet functions [8], which are well explored in mathematical statistics [1, 2]. The second variant of results (i) and (ii) as well as result (iii) apply results from stochastic programming [17].

The rest of this paper is structured as follows: Section 2 constructs the orbit space of partitions and introduces metric structures. In Section 3, we introduce Fréchet consensus functions and study their asymptotic behavior. Section 4 presents examples of continuous extensions of standard cluster criteria. Finally, Section 5 concludes with a summary of the main results and with an outlook to further research. We present proofs in the appendix.

2. Geometry of Partition Spaces

In this section, we show that a partition can be represented as a point in some geometric space, called orbit space. Then we endow orbit spaces $\mathcal{P}$ with metrics δ derived from the Euclidean space and study their properties.

2.1. Partitions

Let $\mathcal{Z} = \{z_1, \ldots, z_m\}$ be a set of m data points. A partition X of $\mathcal{Z}$ with ℓ clusters $\mathcal{C}_1, \ldots, \mathcal{C}_\ell$ is specified by a matrix $\boldsymbol{X} \in [0,1]^{\ell \times m}$ such that $\boldsymbol{X}^T \mathbf{1}_\ell = \mathbf{1}_m$, where $\mathbf{1}_\ell \in \mathbb{R}^\ell$ and $\mathbf{1}_m \in \mathbb{R}^m$ are vectors of all ones.

The rows $\boldsymbol{x}_{k:}$ of matrix $\boldsymbol{X}$ refer to the clusters $\mathcal{C}_k$ of partition X. The columns $\boldsymbol{x}_{:j}$ of $\boldsymbol{X}$ refer to the data points $z_j \in \mathcal{Z}$. The elements x_{kj} of matrix $\boldsymbol{X} = (x_{kj})$ represent the degree of membership of data point z_j to cluster $\mathcal{C}_k$. The constraint $\boldsymbol{X}^T \mathbf{1}_\ell = \mathbf{1}_m$ demands that the membership values $\boldsymbol{x}_{:j}$ of data point z_j across all clusters must sum to one.

By $\mathcal{P}_{\ell,m}$ we denote the set of all partitions with ℓ clusters over m data points. Since some clusters may be empty, the set $\mathcal{P}_{\ell,m}$ also contains partitions with less than ℓ clusters. Thus, we consider $\ell \leq m$ as the maximum number of clusters we encounter. If the exact numbers ℓ and m do not matter or are clear from the context, we also write $\mathcal{P}$ for $\mathcal{P}_{\ell,m}$. A hard partition X is a partition with matrix representation $\boldsymbol{X} \in \{0,1\}^{\ell \times m}$. The set $\mathcal{P}^+ \subset \mathcal{P}$ denotes the subset of all hard partitions.

2.2. The Orbit Space of Partitions

The representation space $\mathcal{X}$ of the set $\mathcal{P} = \mathcal{P}_{\ell,m}$ of partitions is a set of the form

$$\mathcal{X} = \left\{ \boldsymbol{X} \in [0,1]^{\ell \times m} \,:\, \boldsymbol{X}^T \mathbf{1}_\ell = \mathbf{1}_m \right\}.$$

Then we have a natural projection

$$\pi : \mathcal{X} \to \mathcal{P}, \quad \boldsymbol{X} \mapsto X = \pi(\boldsymbol{X})$$

that sends matrices $\boldsymbol{X}$ to partitions X they represent. The map π conveys two properties: (1) each partition can be represented by at least one matrix, and (2) a partition may have several matrix representations.

Suppose that matrix $\boldsymbol{X} \in \mathcal{X}$ represents a partition $X \in \mathcal{P}$. The subset of all matrices representing X forms an equivalence class $[\boldsymbol{X}]$ that can be obtained by permuting the rows of matrix $\boldsymbol{X}$ in all possible ways. The equivalence class of $\boldsymbol{X}$, called orbit henceforth, is of the form

$$[\boldsymbol{X}] = \{\boldsymbol{P}\boldsymbol{X} \,:\, \boldsymbol{P} \in \Pi\},$$

where Π is the group of all $(\ell \times \ell)$-permutation matrices. The orbit space of partitions is the set

$$\mathcal{X}/\Pi = \{[\boldsymbol{X}] \,:\, \boldsymbol{X} \in \mathcal{X}\}.$$

The orbit space consists of all orbits $[\boldsymbol{X}]$, we can construct as described above. Mathematically, the orbit space $\mathcal{X}/\Pi$ is the quotient space obtained by the action of the permutation group Π on the set $\mathcal{X}$. The orbits $[\boldsymbol{X}]$ are in 1-1-correspondence with the partitions $X = \pi(\boldsymbol{X})$. Therefore, we can identify partitions with orbits and occasionally write $\boldsymbol{X} \in X$ if $X = \pi(\boldsymbol{X})$.

2.3. Metric Structures

This section endows the partition space $\mathcal{P}$ with metrics that are derived by a generic construction principle. As examples, we consider metrics δ_p derived from l_p-metrics of Euclidean spaces. We show that $(\mathcal{P}, \delta_p)$ is a compact metric space for $p \geq 1$, and $(\mathcal{P}, \delta_2)$ is a geodesic space.

Every metric d on the representation space $\mathcal{X} \subset \mathbb{R}^{\ell \times m}$ induces a distance function

$$\delta : \mathcal{P} \times \mathcal{P} \to \mathbb{R}, \quad (X, Y) \mapsto \min\{d(\boldsymbol{X}, \boldsymbol{Y}) \,:\, \boldsymbol{X} \in X, \boldsymbol{Y} \in Y\}. \tag{1}$$

Note that the minimum in (1) exists, because the orbits $[\boldsymbol{X}]$ and $[\boldsymbol{Y}]$ are finite. As an example, we consider distance functions induced by the l_p-norm. The l_p-norm for matrices $\boldsymbol{X} \in \mathcal{X}$ is defined by

$$\|\boldsymbol{X}\|_p = \left(\sum_{k=1}^{\ell} \sum_{j=1}^{m} |x_{kj}|^p \right)^{1/p}$$

for every $p \geq 1$. The l_p-norm induces the distance function

$$\delta_p : \mathcal{P} \times \mathcal{P} \to \mathbb{R}, \quad (X, Y) \mapsto \min\left\{ \|\boldsymbol{X} - \boldsymbol{Y}\|_p \,:\, \boldsymbol{X} \in X, \boldsymbol{Y} \in Y \right\},$$

called l_p-distance on $\mathcal{P}$, henceforth.

To show that distances δ on $\mathcal{P}$ induced by metrics d on $\mathcal{X}$ are also metrics, we demand that metric d is permutation invariant. We say, a metric d on $\mathcal{X}$ is permutation invariant, if

$$d(\boldsymbol{P}\boldsymbol{X}, \boldsymbol{P}\boldsymbol{Y}) = d(\boldsymbol{X}, \boldsymbol{Y})$$

for all permutations $\boldsymbol{P} \in \Pi$. Permutation invariance means that the metric d is invariant under simultaneously relabeling the clusters of $\boldsymbol{X}$ and $\boldsymbol{Y}$. An example of permutation invariant metrics are the l_p-metrics. The next result shows that permutation invariant metrics on $\mathcal{X}$ induce distances on $\mathcal{P}$ that are again metrics.

Theorem 2.1. *Let $(\mathcal{X}, d)$ be a metric space and let $(\mathcal{P}, \delta)$ be the partition space endowed with distance function δ induced by metric d. Suppose that d is permutation invariant. Then we have:*

1. *The distance δ is a metric.*
2. *$(\mathcal{P}, \delta)$ is a compact space.*
3. *$(\mathcal{P}, \delta_2)$ is a geodesic space.*

Theorem 2.1 presents a generic way to construct metrics on $\mathcal{P}$. Being a compact metric space is a strong property for consistency statements. Being a geodesic space means that any pair of partitions X and Y have a midpoint partition M such that

$$\delta_2(X, M) = \delta_2(Y, M) = \frac{1}{2}\delta_2(X, Y).$$

Note that being a geodesic space is a necessary and sufficient condition for guaranteeing the midpoint property for all pairs of partitions.

3. Fréchet Consensus Clustering

This section first formalizes the problem of consensus clustering using the mean partition approach and then studies its asymptotic behavior. For this, we link the consensus function of the mean partition approach to the Fréchet function [8] from mathematical statistics. Then we show that under normal conditions the mean partition approach is consistent and asymptotically normal.

3.1. Fréchet Functions

Let $\rho : \mathcal{P} \times \mathcal{P} \to \mathbb{R}$ be a (dis)similarity function on the partition space $\mathcal{P}$. To avoid case distinctions between min- and max-operations, we assume that $\rho = -s$ if s is a similarity function. Typically, we may think of ρ as a criterium for comparing partitions such as, for example, the metrics δ_p, the Mirkin distance, the Rand index, and the variation of information [15].

We assume that Q is a probability distribution on the partition space $\mathcal{P}$. Suppose that $\mathcal{S}_n = (X_1, X_2, \ldots, X_n)$ is a sample of n partitions $X_i \in \mathcal{P}$ drawn i.i.d. from the probability distribution Q. Then

$$F_n : \mathcal{P} \to \mathbb{R}, \quad Z \mapsto \frac{1}{n}\sum_{i=1}^{n} \rho(X_i, Z) \tag{2}$$

is the Fréchet function of $\mathcal{S}_n$ corresponding to the dissimilarity function ρ. The infimum V_n of F_n is the variation of $\mathcal{S}_n$. The minimizers of the Fréchet function need neither exist nor be unique. The (possibly empty) set

$$\mathcal{F}_n = \left\{ M \in \mathcal{P} \,:\, M = \arg\min_Z F_n(Z) \right\}$$

is called the mean partition set of $\mathcal{S}_n$. Every $M \in \mathcal{F}_n$ is a mean partition of $\mathcal{S}_n$.

In statistics, the mean is an estimator of the expectation. To adopt this idea to partitions, we define the expected Fréchet function of Q corresponding to the dissimilarity function ρ as

$$F_Q : \mathcal{P} \to \mathbb{R}, \quad Z \mapsto \int_{\mathcal{P}} \rho(X, Z)\, dQ(X).$$

The infimum V_Q of F_Q is the expected variation of Q. The set

$$\mathcal{F}_Q = \left\{ M \in \mathcal{P} \,:\, M = \arg\min_Z F_Q(Z) \right\}$$

is called the expected partition set of Q. Any $M \in \mathcal{F}_Q$ is an expected partition of Q.

3.2. Consistency by Metric Structures

This section shows consistency of the mean partition set $\mathcal{F}_n$ and the variation V_n under the assumption that the dissimilarity function ρ is of the form

$$\rho = h \circ \delta,$$

where δ is a metric on $\mathcal{P}$ and $h : \mathbb{R}_+ \to \mathbb{R}_+$ is a continuous non-negative loss function. The consistency result of this section is the first of two variants. The distinguishing feature of this variant is that consistency of mean partitions is shown in an abstract metric space without connection to an ambient Euclidean space.

Roughly, we may think of the mean partition set $\mathcal{F}_n$ being a strongly consistent estimator of the expected partition set $\mathcal{F}_Q$, if

$$P\left(\lim_{n\to\infty} \mathcal{F}_n \subseteq \mathcal{F}_Q \right) = 1.$$

A precise definition of strong consistency in the sense of Bhattacharya-Patrangenaru (BP) is given in Section A. We have the following existence and consistency result:

Theorem 3.1. *Suppose that $(\mathcal{P}, \delta)$ is a compact metric space and $h : \mathbb{R}_+ \to \mathbb{R}_+$ is a continuous loss function. Then the following holds:*

1. *$\mathcal{F}_Q$ are non-empty and compact.*
2. *$\mathcal{F}_n$ is a BP-strongly consistent estimator of $\mathcal{F}_Q$.*
3. *If $\mathcal{F}_Q = \{M\}$, then $\mathcal{F}_n$ is a strongly consistent estimator of M.*

4. V_n is a strongly consistent estimator of V_Q.

Proof. The assertions of Theorem 3.1 follow from [2], Theorem 2.3, Prop. 2.8, and Corollary 2.4. ■

The role of the loss function h is to generalizes the concept of Fréchet function of order p corresponding to $\rho(X,Y) = \delta(X,Y)^p$. In this case, the loss is of the form $h(x) = x^p$. For $p = 1$, we recover the consensus function of the median partition and for $p = 2$ the consensus function of the mean partition.

Immediate consequences of Theorem 2.1 and 3.1 are as follows:

Corollary 3.2. *Let $(\mathcal{P}, \delta)$ be a partition space. Assertions (1)–(4) of Theorem 3.1 hold for the following special cases:*

1. *δ is induced by a permutation invariant metric on $\mathcal{X}$.*
2. *$\delta = \delta_p$ is an l_p-metric on $\mathcal{P}$.*

3.3. Consistency by Stochastic Programming

In this section, we consider the following stochastic programming problem:

$$\min_{Z \in \mathcal{P}} \left\{ F_Q(Z) = \int_{\mathcal{P}} \rho(X,Z)\, dQ(X) \right\}, \tag{3}$$

where $\rho(X,Z)$ is a continuous dissimilarity function and X is a random partition with probability distribution Q. The consistency result of this section is the second of both variants. This variant requires no metric structure of the dissimilarity function. However, the proof compensates the missing structure of ρ by exploiting the structure of the ambient Euclidean space.

To establish consistency, we need to introduce two concepts: (1) continuity of a function on partitions, and (2) distance between partitions.

To define both concepts, we assume the geodesic space $(\mathcal{P}, \delta_2)$, where δ_2 is the l_2-metric induced by the Euclidean norm. We say, a function $f : \mathcal{P} \to \mathbb{R}$ is continuous at partition Z, if for every $\varepsilon > 0$ there is a $\zeta > 0$ such that for all $X \in \mathcal{P}$ with $\delta_2(X,Z) < \zeta$ the value $f(X)$ satisfies $|f(X) - f(Z)| < \varepsilon$. Furthermore, by

$$D(\mathcal{U}, \mathcal{V}) = \sup_{X \in \mathcal{U}} \inf_{Y \in \mathcal{V}} \delta_2(X,Y)$$

we denote the distance between two subsets $\mathcal{U}, \mathcal{V} \subseteq \mathcal{P}$.

The next theorem gives convergence results based on the continuity of the dissimilarity function ρ.

Theorem 3.3. *Let $X_1, \ldots, X_n \in \mathcal{P}$ be a sample of n partitions drawn i.i.d. from the probability distribution Q. Suppose that $\rho(X,Z)$ is continuous in both arguments. Then*

1. $D(\mathcal{F}_n, \mathcal{F}_Q) \to 0$ *almost surely.*

2. $V_n \to V_Q$ *almost surely.*

Section 4 presents examples of continuous dissimilarity functions.The examples are continuos extensions of standard criteria for comparing partitions. For all these examples, the mean partition approach is consistent in the sense of Theorem 3.3 under mild assumptions.

3.4. A Central Limit Theorem

Consistency of the mean partition approach gives a certain assurance that the error of the estimation tends to zero in the limit as the sample size grows to infinity. It does, however, not indicate the magnitude of the error for a given sample size. In standard statistics, we can use the Central Limit Theorem (CLT) to construct an approximate confidence interval for the unknown population mean and derive the order of the estimation error as a function of the sample size. To lay the foundations for these techniques, we present a variant of the CLT for mean partitions.

We assume the same setting as in Section 3.3. By $\mathcal{N}(0, \sigma^2)$ we denote the normal distribution with mean zero and variance σ^2. Convergence of random variable (X_n) to a random variable X in distribution is denoted by $X_n \xrightarrow{d} X$. The next result states a version of the CLT for mean partitions.

Theorem 3.4. *Suppose that every sample $\mathcal{S}_n$ of n partitions is drawn i.i.d. from $\mathcal{P}$. Then*

$$\frac{V_n - V}{\sqrt{n}} \xrightarrow{d} \inf_{M \in \mathcal{F}_Q} \mathcal{N}(0, \sigma^2(M)).$$

If the expected partition set $\mathcal{F}_Q = \{M\}$ consists of a singleton M, then Theorem 3.4 reduces to

$$\frac{V_n - V}{\sqrt{n}} \xrightarrow{d} \mathcal{N}(0, \sigma^2(M)).$$

4. Generalized Criteria for Comparing Partitions

This section extends standard criteria for comparing hard partitions to criteria for comparing arbitrary partitions. There are infinitely many ways to extend functions from a discrete to a continuous domain. We suggest extensions that partially admit a probabilistic interpretation. Continuity of a criterium allows us to invoke Theorem 3.3 and 3.4. For details on criteria based on counting pairs and cluster matchings, we refer to [15] and for details on information-theoretic criteria, we refer to [23].

4.1. Criteria based on Counting Pairs

A hard partition $X \in \mathcal{P}^+$ induces an equivalence relation

$$z \sim_X z' \quad \Leftrightarrow \quad z \text{ and } z' \text{ are in the same cluster of } X$$

for all $z, z' \in \mathcal{Z}$. Let $\mathcal{Z}^{[2]}$ denote the set of all 2-element subsets of $\mathcal{Z}$. The confusion matrix $\boldsymbol{C}^+(X,Y) = (m^+_{pq})$ of hard partitions X and Y is a (2×2)-matrix with elements

$$\begin{aligned}
m^+_{11} &= \left|\left\{\{z,z'\} \in \mathcal{Z}^{[2]} \,:\, z \sim_X z', z \sim_Y z'\right\}\right| \\
m^+_{10} &= \left|\left\{\{z,z'\} \in \mathcal{Z}^{[2]} \,:\, z \sim_X z', z \nsim_Y z'\right\}\right| \\
m^+_{01} &= \left|\left\{\{z,z'\} \in \mathcal{Z}^{[2]} \,:\, z \nsim_X z', z \sim_Y z'\right\}\right| \\
m^+_{00} &= \left|\left\{\{z,z'\} \in \mathcal{Z}^{[2]} \,:\, z \nsim_X z', z \nsim_Y z'\right\}\right|.
\end{aligned}$$

The matrix $\boldsymbol{C}^+(X,Y)$ satisfies

$$\left\|\boldsymbol{C}^+(X,Y)\right\|_1 = m^+_{11} + m^+_{10} + m^+_{01} + m^+_{00} = \frac{m(m-1)}{2} = \left|\mathcal{Z}^{[2]}\right|.$$

By $N = m(m-1)/2$ we denote the cardinality of the set $\mathcal{Z}^{[2]}$. Criteria based on counting pairs can be described by a function of the general form

$$\rho(X,Y) = f\Big(\boldsymbol{C}^+(X,Y)\Big),$$

where $X, Y \in \mathcal{P}^+$. To extend ρ to the space $\mathcal{P}$ of all partitions, we define the compatibility matrix of a partition $X \in \mathcal{P}$ as an $(m \times m)$-matrix of the form

$$\boldsymbol{C}_X = \boldsymbol{X}^T\boldsymbol{X},$$

where $\boldsymbol{X}$ is an arbitrary representation of X. From Prop. B.4 follows that $\boldsymbol{C}_X$ is independent from the particular choice of representation $\boldsymbol{X} \in X$. The compatibility matrix $\boldsymbol{C}_X = (c_{rs})$ consists of elements of the form

$$c_{rs} = \langle \boldsymbol{x}_{:r}, \boldsymbol{x}_{:s} \rangle$$

for all $r, s \in \{1, \ldots, m\}$. Thus, c_{rs} is the inner product of columns $\boldsymbol{x}_{:r}$ and $\boldsymbol{x}_{:s}$ of matrix $\boldsymbol{X}$. Recall that a column $\boldsymbol{x}_{:j}$ of a representation matrix $\boldsymbol{X}$ summarizes the membership values of data point $z_j \in \mathcal{Z}$.

Next, we define the function

$$\chi(\boldsymbol{A}, \boldsymbol{B}) = \sum_{r=1}^{m} \sum_{s=r+1}^{m} a_{rs} b_{rs}$$

for all $(m \times m)$-matrices $\boldsymbol{A} = (a_{rs})$ and $\boldsymbol{B} = (b_{rs})$. The function $\chi(\boldsymbol{A}, \boldsymbol{B})$ is the inner product of the strictly upper triangular matrices of $\boldsymbol{A}$ and $\boldsymbol{B}$. By $\mathbf{1} = \mathbf{1}_{m\times m}$ we denote the $(m \times m)$-matrix of all ones.

Now we can extend the confusion matrix $\boldsymbol{C}^+(X,Y)$ to its continuous counterpart $\boldsymbol{C}(X,Y) = (m_{pq})$ with elements

$$\begin{aligned} m_{11} &= \chi\left(\boldsymbol{C}_X, \boldsymbol{C}_Y\right) \\ m_{10} &= \chi\left(\boldsymbol{C}_X, \mathbf{1} - \boldsymbol{C}_Y\right) \\ m_{01} &= \chi\left(\mathbf{1} - \boldsymbol{C}_X, \boldsymbol{C}_Y\right) \\ m_{00} &= \chi\left(\mathbf{1} - \boldsymbol{C}_X, \mathbf{1} - \boldsymbol{C}_Y\right), \end{aligned}$$

The next result shows that matrix $\boldsymbol{C}(X,Y)$ indeed generalizes the confusion matrix $\boldsymbol{C}^+(X,Y)$ and admits a probabilistic interpretation.

Proposition 4.1. *Let $X, Y \in \mathcal{P}$ be partitions of dataset $\mathcal{Z}$ consisting of m elements. The matrix $\boldsymbol{C}(X,Y) = (m_{pq})$ satisfies the following properties:*

1. *If $X, Y \in \mathcal{P}^+$, then $\boldsymbol{C}^+(X,Y) = \boldsymbol{C}(X,Y)$.*
2. *$m_{pq} \geq 0$ for all $p, q \in \{0,1\}$.*
3. *$\sum_{p,q} m_{pq} = N$.*

The first assertion of Prop. 4.1 states that $\boldsymbol{C}(X,Y)$ is an extension of the confusion matrix $\boldsymbol{C}^+(X,Y)$. The second and third assertion admit a probabilistic interpretation of the values m_{pq} as follows:

$$\begin{aligned} \mathbb{P}(z \sim_X z', z \sim_Y z') &= \frac{m_{11}}{\left|\mathcal{Z}^{[2]}\right|} \\ \mathbb{P}(z \sim_X z', z \nsim_Y z') &= \frac{m_{10}}{\left|\mathcal{Z}^{[2]}\right|} \\ \mathbb{P}(z \nsim_X z', z \sim_Y z') &= \frac{m_{01}}{\left|\mathcal{Z}^{[2]}\right|} \\ \mathbb{P}(z \nsim_X z', z \nsim_Y z') &= \frac{m_{00}}{\left|\mathcal{Z}^{[2]}\right|}, \end{aligned}$$

where $\{z, z'\} \in \mathcal{Z}^{[2]}$. The probabilistic interpretation of m_{pq} is consistent with the probabilistic interpretation of m_{pq}^+ for hard partitions.

We can extend criteria $\rho(X,Y) = f(\boldsymbol{C}^+(X,Y))$ based on counting pairs from the space $\mathcal{P}^+$ of hard partitions to the space $\mathcal{P}$ of all partitions by replacing the confusion matrix $\boldsymbol{C}^+(X,Y)$ with matrix $\boldsymbol{C}(X,Y)$. We present extensions of some common

examples:

$$\rho_1(X,Y) = \frac{m_{11}}{m_{11}+m_{01}} \qquad \text{(Wallace I)}$$

$$\rho_2(X,Y) = \frac{m_{11}}{m_{11}+m_{01}} \qquad \text{(Wallace II)}$$

$$\rho_3(X,Y) = \frac{m_{11}+m_{00}}{N} \qquad \text{(Rand)}$$

$$\rho_4(X,Y) = \frac{m_{11}}{\sqrt{(m_{11}+m_{10})\,(m_{11}+m_{01})}} \qquad \text{(Fowlkes-Mallows)}$$

$$\rho_5(X,Y) = \frac{m_{11}}{m_{11}+m_{01}+m_{01}} \qquad \text{(Jacard)}$$

The next result shows that criterium ρ_3 is continuous and criteria ρ_1 – ρ_5 are continuous almost everywhere.

Proposition 4.2. *Suppose that $m > 1$. Then criteria $\rho_1 - \rho_5$ are continuous in both arguments almost everywhere. Criterium ρ_3 is continuous.*

Critical points occur when we admit partitions for which $m_{11} = 0$. This situation refers to the degenerated case where each cluster contains at most one data point with positive membership value. The simplest way to cope with this issue is to consider closed subsets $\mathcal{P}'$ of $\mathcal{P}$ that exclude degenerated partitions. Then Theorem 3.3 and 3.4 can be applied to ρ_1 – ρ_5 restricted to the compact subset $\mathcal{P}'$.

4.2. Criteria based on Cluster Matchings

Every hard partition X defines ℓ clusters $\mathcal{C}_1(X), \ldots, \mathcal{C}_\ell(X)$ that partition the data set $\mathcal{Z}$. For two hard partitions X and Y, we define the numbers

$$\begin{aligned} m &= |\mathcal{Z}| \\ x_p^+ &= |\mathcal{C}_p(X)| \\ y_q^+ &= |\mathcal{C}_q(Y)| \\ z_{pq}^+ &= |\mathcal{C}_p(X) \cap \mathcal{C}_q(Y)| \end{aligned}$$

for all $p, q \in \{1, \ldots, \ell\}$. Criteria based on cluster matchings can be described by a function of the general form

$$\rho(X,Y) = f\Big(m, \big(x_p^+\big), \big(y_q^+\big), \big(z_{pq}^+\big)\Big),$$

where $X, Y \in \mathcal{P}^+$. Let $\boldsymbol{X}$ and $\boldsymbol{Y}$ be representations of partitions X and Y from $\mathcal{P}$. To extend a criterium ρ to the space of all partitions, we define the values

$$\begin{aligned} x_p &= \langle \boldsymbol{x}_{p:}, \mathbf{1}_m \rangle \\ y_q &= \langle \boldsymbol{y}_{q:}, \mathbf{1}_m \rangle \\ z_{pq} &= \langle \boldsymbol{x}_{p:}, \boldsymbol{y}_{q:} \rangle \end{aligned}$$

for all $p, q \in \{1, \ldots, \ell\}$. Recall that the rows $\boldsymbol{x}_{p:}$ and $\boldsymbol{y}_{q:}$ of $\boldsymbol{X}$ and $\boldsymbol{Y}$represent the clusters $\mathcal{C}_p(X)$ and $\mathcal{C}_q(Y)$.

The next result shows that (x_p), (y_p), (z_{pq}) generalize (x_p^+), (y_p^+), (z_{pq}^+) and partially admit a probabilistic interpretation.

Proposition 4.3. *Let $X, Y \in \mathcal{P}$ be partitions of dataset $\mathcal{Z}$ consisting of m elements. Then we have:*

1. *If $X, Y \in \mathcal{P}^+$, then $x_p^+ = x_p$, $y_q^+ = y_q$, and $z_{pq}^+ = z_{pq}$.*
2. *$x_p, y_q, z_{pq} \geq 0$ for all $p, q \in \{1, \ldots, \ell\}$.*
3. *$\sum_p x_p = \sum_q y_q = m$*

The first assertion of Prop. 4.1 states that (x_p), (y_p), (z_{pq}) are extensions. The second and third assertion admit a probabilistic interpretation as follows:

$$\mathbb{P}(z \in \mathcal{C}_p(X)) = \frac{x_p}{|\mathcal{Z}|}$$
$$\mathbb{P}(z \in \mathcal{C}_q(Y)) = \frac{y_q}{|\mathcal{Z}|}$$

Next, we extend criteria $\rho(X, Y) = f\left(m, \left(x_p^+\right), \left(y_q^+\right), \left(z_{pq}^+\right)\right)$ based on cluster matchings from the space $\mathcal{P}^+$ of hard partitions to the space $\mathcal{P}$:

$$\rho_6(X, Y) = \sum_p x_p^2 + \sum_q y_q^2 - \sum_p \sum_q z_{pq}^2 \qquad \text{(Mirkin)}$$

$$\rho_7(X, Y) = \max_{\phi \in \mathrm{Sym}_\ell} \frac{1}{m} \sum_p z_{p\phi(q)} \qquad \text{(Meilă-Heckerman)}$$

$$\rho_8(X, Y) = 2m - \sum_p \max_q z_{pq} - \sum_q \max_p z_{pq} \qquad \text{(van Dongen)}$$

The set Sym_ℓ in the Meilă-Heckerman formula is the set of all permutations of the numbers $1, \ldots, \ell$.

Since criteria ρ_6–ρ_8 are continuous, the mean partitions are consistent estimators of the expected partitions according to Theorem 3.3 and asymptotic normal according to Theorem 3.4.

4.3. Criteria based on Information Theory

Suppose that X and Y are hard partitions. Similarly as criteria based on cluster matchings, criteria based on Information Theory can be described by a function of the general form

$$\rho(X, Y) = f\Big(m, \left(x_p^+\right), \left(y_q^+\right), \left(z_{pq}^+\right)\Big),$$

where $X, Y \in \mathcal{P}^+$. In contrast to criteria based on cluster matchings, the function f is composed of information-theoretic measures. Using the same notation as in the

previous sections, continuous extensions of information-theoretic criteria for partitions $X, Y \in \mathcal{P}$ are based on the following measures:

$$H(X) = -\sum_p \frac{x_p}{N} \log \frac{x_p}{N} \qquad \text{(entropy)}$$

$$H(X,Y) = -\sum_p \sum_q \frac{z_{pq}}{N} \log \frac{z_{pq}}{N} \qquad \text{(joint entropy)}$$

$$H(X|Y) = -\sum_p \sum_q \frac{z_{pq}}{N} \log \frac{z_{pq}/N}{y_q/N} \qquad \text{(conditional entropy)}$$

$$I(X,Y) = \sum_p \sum_q \frac{z_{pq}}{N} \log \frac{z_{pq}/N}{x_p y_q/N^2}. \qquad \text{(mutual information)}$$

Recall that $N = m(m-1)/2$. Continuous extensions of information-theoretic cluster criteria take the general form

$$\rho(X,Y) = f\Big(H(X), H(X,Y), H(X|Y), I(X,Y)\Big)$$

for all $X, Y \in \mathcal{P}$. For examples of information-theoretic measures, we refer to [23].

By considering closed subsets of $\mathcal{P}$ that exclude degenerated partitions, we can invoke Theorem 3.3 and 3.4 for function f that are continuous almost everywhere.

5. Conclusion

Under normal assumptions, consensus clustering based on the mean partition approach is consistent and asymptotic normal. To derive these results, we represented partitions as points of an orbit space and established links to the theory of Fréchet means and stochastic programming. The results hold for continuous extensions of standard cluster criteria under mild assumptions.

The orbit space approach provides a mathematical foundation for studying further geometrical, analytical, and statistical properties of sets of partitions with applications not confined to clustering but for any domain that studies equivalence relations. The theoretical results suggest to apply algorithms from stochastic optimization for minimizing Fréchet consensus functions.

A. Bhattacharya-Patrangenaru Consistency

The metric δ of the partition space $(\mathcal{P}, \delta)$ induces a Borel σ-algebra $\mathcal{B}$ on $\mathcal{P}$ such that $(\mathcal{P}, \mathcal{B})$ is a measurable space. Let $(\Omega, \mathcal{A}, P)$ be an abstract probability space. We assume that $(\Omega, \mathcal{A}, P)$ is complete in the sense that every subset of every null-set is measurable.

A random element X taking values in $\mathcal{P}$ is a mapping $X : \Omega \to \mathcal{P}$ measurable with respect to the Borel σ-algebra $\mathcal{B}$. Then $Q = P \circ X^{-1}$ is a probability measure on the measurable space $(\mathcal{P}, \mathcal{B})$.

The sample mean set $\mathcal{F}_n$ is a Bhattacharya-Patrangenaru (BP) strongly consistent estimator of the mean set $\mathcal{F}_Q$, if $\mathcal{F}_Q \neq \emptyset$ and if for every $\varepsilon > 0$ and for almost every $\omega \in \Omega$ there is an integer $N = N(\varepsilon, \omega)$ such that

$$\mathcal{F}_n \subseteq \mathcal{F}_Q^{\varepsilon} = \{Z \in \mathcal{P} \,:\, \delta\,(Z, \mathcal{F}_Q) \leq \varepsilon\}$$

for all $n \geq N$ [1, 2].

B. Proofs

B.1. Proof of Theorem 2.1

As preparation for the third assertion of Theorem 2.1, we first show that partitions spaces $(\mathcal{P}_{\ell,m}, \delta_2)$ can be isometrically embedded into graph spaces endowed with a graph edit kernel metric as defined in [11].

Let $\mathcal{A} = \mathbb{R}^m$ be the set of node and edge attributes. An attributed graph of order ℓ is a triple $X = (\mathcal{V}, \mathcal{E}, \alpha)$, where $\mathcal{V}$ represents a set of ℓ nodes, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ a set of edges, and $\alpha : \mathcal{V} \times \mathcal{V} \to \mathcal{A}$ is an attribute function satisfying the following properties for all $i, j \in \mathcal{V}$ with $i \neq j$:

1. $\alpha(i, j) \neq \mathbf{0}$ for all edges $(i, j) \in \mathcal{E}$
2. $\alpha(i, j) = \mathbf{0}$ for all non-edges $(i, j) \notin \mathcal{E}$.

As opposed to edges, node attributes $\alpha(i, i)$ may take any value from $\mathcal{A}$. By $\mathcal{G}_{\ell,m}$ we denote the set of all attributed graphs of order ℓ with attributes from $\mathcal{A} = \mathbb{R}^m$. The graph edit kernel metric is a function of the from

$$\delta_g : \mathcal{G}_{\ell,m} \times \mathcal{G}_{\ell} \to \mathbb{R}, \quad (X, Y) \mapsto \min_{\phi \in \mathrm{Sym}_{\ell}} \sum_{i,j} \|\alpha(i, j) - \alpha(\phi(i), \phi(j))\|^2,$$

where Sym_{ℓ} is the symmetric group of permutations that can be performed on a set of cardinality ℓ. The pair $(\mathcal{G}_{\ell,m}, \delta_g)$ is called graph edit kernel space of order ℓ.

Theorem B.1. *There is an isometric embedding of the orbit space $(\mathcal{P}_{\ell,m}, \delta_2)$ into the graph edit kernel space $(\mathcal{G}_{\ell,m}, \delta_g)$.*

Proof. Let $X \in \mathcal{P} = \mathcal{P}_{\ell,m}$ be a partition with matrix representation $\boldsymbol{X}$. A partition X can be equivalently expressed by a graph $G_X = (\mathcal{V}, \emptyset, \alpha)$ as follows: The set $\mathcal{V}$ consists of ℓ nodes, each of which represents a cluster. The set of edges is empty. The attribute function is of the form

$$\alpha(k, j) = \begin{cases} \boldsymbol{x}_{k:} & : \quad k = j \\ \mathbf{0} & : \quad k \neq j \end{cases},$$

where $\boldsymbol{x}_{k:}$ is the k-th row of matrix $\boldsymbol{X}$ referring to the k-th cluster of X. We call G_X the partition graph of X. This construction gives an injective map $f : \mathcal{P}_{\ell,m} \to \mathcal{G}_{\ell,m}$.

Next, we show that the map f is isometric. Observe that a graph $G = (\mathcal{V}, \mathcal{E}, \alpha) \in \mathcal{G}_{\ell,m}$ can be represented by a $(\ell \times \ell)$-matrix $\boldsymbol{G} = (\boldsymbol{g}_{ij})$ with elements $\boldsymbol{g}_{ij} = \alpha(i, j)$. Then the graph edit kernel distance between graphs G and H with matrix representations $\boldsymbol{G}$ and $\boldsymbol{H}$, resp., can be equivalently expressed by

$$\delta_g(G, H) = \min_{\boldsymbol{P} \in \Pi} \left\| \boldsymbol{G} - \boldsymbol{P}\boldsymbol{H}\boldsymbol{P}^T \right\|_2,$$

where Π is the set of all $(\ell \times \ell)$-permutation matrices [11]. Let $X, Y \in \mathcal{P}_{\ell,m}$ be two partitions with respective matrix representations $\boldsymbol{X}$ and $\boldsymbol{Y}$. Suppose that G_X and G_Y are the partition graphs of X and Y, resp., with respective matrix representations $\boldsymbol{G}_X$ and $\boldsymbol{G}_Y$. Since $\boldsymbol{G}_X$ and $\boldsymbol{G}_Y$ are diagonal matrices, we have

$$\|\boldsymbol{X} - \boldsymbol{P}\boldsymbol{Y}\|_2 = \left\|\boldsymbol{G}_X - \boldsymbol{P}\,\boldsymbol{G}_Y\boldsymbol{P}^T\right\|_2$$

for all $\boldsymbol{P} \in \Pi$. This shows

$$\delta_2(X,Y) = \min_{\boldsymbol{P}\in\Pi} \|\boldsymbol{X} - \boldsymbol{P}\boldsymbol{Y}\|_2 = \min_{\boldsymbol{P}\in\Pi} \left\|\boldsymbol{G}_X - \boldsymbol{P}\,\boldsymbol{G}_Y\boldsymbol{P}^T\right\|_2 = \delta_g\left(\boldsymbol{G}_X, \boldsymbol{G}_Y\right).$$

Thus, f is an isometric embedding. ■

Proof of Theorem 2.1. Note that Π is a group of isometries acting on $\mathcal{X}$, because metric d is invariant under permutations by assumption.

1. Since Π is finite, each orbit $[\boldsymbol{X}]$ is a finite and therefore a closed subset of $\mathcal{X}$. Then the distance function δ induced by metric d is a metric by [16], Theorem 6.6.1.

2. By construction of δ, the group Π is a finite group of isometries acting on the space $\mathbb{R}^{\ell\times m}$. Then all orbits $[\boldsymbol{X}]$ are finite and closed subsets of $\mathbb{R}^{\ell\times m}$. Since the Euclidean space is a finitely compact metric space, the quotient space $\mathbb{R}^{\ell\times m}/\Pi$ is a complete metric space by [16], Theorem 8.5.2. Since $\mathcal{X} \subseteq \mathbb{R}^{\ell\times m}$ is compact and Π is finite and therefore a compact group, compactness of $\mathcal{P}$ follows from [3], Chapter I, Theorem 3.1.

3. We show that $(\mathcal{P}_{\ell,m}, \delta_2)$ is a geodesic space. From [11], Theorem 3.3 follows that $(\mathcal{G}_{\ell,m}, \delta_g)$ is a geodesic space. Let $f : \mathcal{P}_{\ell,m} \to \mathcal{G}_{\ell,m}$ be the isometric embedding defined in the proof of Theorem B.1. Suppose that X and Y are partitions from $\mathcal{P}_{\ell,m}$ and $G_X = f(X)$ and $G_Y = f(Y)$ the corresponding graph partitions. Since $(\mathcal{G}_{\ell,m}, \delta_g)$ is a geodesic space there is a midpoint G_M of G_X and G_Y. From [12], Corollary 3.6 follows that G_M is a minimum of the sample Fréchet function $F_2(Z) = \delta_g(G_X, Z) + \delta_g(G_Y, Z)$. Then by [12], Theorem 3.1 the matrix representation $\boldsymbol{G}_M$ is of the form

$$\boldsymbol{G}_M = \frac{1}{2}\left(\boldsymbol{G}_X + \boldsymbol{G}_Y\right).$$

This implies that $G_M \in f(\mathcal{P}_{\ell,m})$. Since G_M is a midpoint and by isometry of f, we find that M with $f(M) = G_M$ is a midpoint of X and Y. From [13], Lemma 2.2 follows that $\mathcal{P}_{\ell,m}$ is a geodesic space.

B.2. Proof of Theorem 3.3

Standard solutions to stochastic programming problems usually assume problem formulations in Euclidean spaces. To invoke these solutions, we first send the stochastic problem (3) to an equivalent problem in the Euclidean space, then solve the problem, and finally project the solution back to the partition space.

Step 1: Moving Problem to Euclidean Space

To transform problem (3) to an equivalent problem in the Euclidean space, we define the pullback of function $\rho(X, Z)$ as

$$\rho^* : \mathcal{X} \times \mathcal{X} \to \mathbb{R}, \quad (\boldsymbol{X}, \boldsymbol{Z}) \mapsto \rho\Big(\pi(\boldsymbol{X}), \pi(\boldsymbol{Z})\Big),$$

where $\pi : \mathcal{X} \to \mathcal{P}$ is the natural projection. The pullback ρ^* of the dissimilarity ρ induces pullbacks of the Fréchet function F_n and the expectation F_Q. The pullback of F_n is of the form

$$F_n^*(\boldsymbol{Z}) = \frac{1}{n} \sum_{i=1}^{n} \rho^*(\boldsymbol{X}_i, \boldsymbol{Z}) \,.$$

By $\mathcal{F}_n^*$ we denote the set of minimizers of F_n^* and by V_n^* the minimum value of F_n^*. Similarly,

$$F_q^*(\boldsymbol{Z}) = \int_{\mathcal{X}} \rho^* \left(\boldsymbol{X}_i, \boldsymbol{Z}\right) dq(\boldsymbol{X})$$

is the pullback of F_Q, where q is a probability measure on $\mathcal{X}$ that induces the probability measure $\mathcal{Q}$ on $\mathcal{P}$ as quotient measure. By $\mathcal{F}_q^*$ we denote the set of minimizers of F_q^* and by V_q^* the minimum value of F_q^*.

Finally, we define the distance between the mean and expected solution set by

$$D^*\left(\mathcal{F}_n^*, \mathcal{F}_q^*\right) = \sup_{\boldsymbol{M}_n \in \mathcal{F}_n^*} \inf_{\boldsymbol{M} \in \mathcal{F}_q^*} \left\| \boldsymbol{M}_n - \boldsymbol{M} \right\|_2 .$$

Step 2: Solving the Problem

By the nature of stochastic programming, we assume that all samples $\mathcal{S}_n$ are drawn i.i.d. We will apply the following Theorem:

[17], Theorem 5.3. *Suppose that there exists a compact set $\mathcal{K} \subset \mathbb{R}^N$ such that*

1. *$\mathcal{F}_q^* \neq \emptyset$ and $\mathcal{F}_q^* \subseteq \mathcal{K}$.*
2. *$F_q^*(\boldsymbol{Z}) < \infty$ is continuous on $\mathcal{K}$.*
3. *$F_n^*(\boldsymbol{Z}) \to F_q^*(\boldsymbol{Z})$ almost surely and uniformly in $\boldsymbol{Z} \in \mathcal{K}$.*
4. *$\mathcal{F}_n^* \neq \emptyset$ and $\mathcal{F}_n^* \subseteq \mathcal{K}$ with probability one.*

Then

1. *$D^*\left(\mathcal{F}_n^*, \mathcal{F}_q^*\right) \to 0$ almost surely.*
2. *$V_n^* \to V_q^*$ almost surely.*

We show that the assumptions of [17], Theorem 5.3 are satisfied. Since $\mathcal{X}$ is compact, we can set $\mathcal{K} = \mathcal{X}$ and $N = \ell \cdot m$.

1. The pullback ρ^* is continuous in both arguments by Lemma B.3. From continuity of ρ^* and compactness of $\mathcal{X}$ follows that F_q^* is continuous. From continuity of F_q^* and again from compactness of $\mathcal{X}$ follows that F_q^* attains it minimum value. Hence, the set $\mathcal{F}_q^* \subseteq \mathcal{X}$ is non-empty and contained in a compact set. This proves the first assumption.

2. To prove the second and third assumption, we invoke the following Theorem:

[17], Theorem 7.48. *Let $\mathcal{S}_n = (\boldsymbol{X}_1, \ldots, \boldsymbol{X}_n)$ be a sample of n partitions $\boldsymbol{X}_i \in \mathcal{X}$ drawn i.i.d. Suppose that:*

1. For any $\boldsymbol{Z} \in \mathcal{K}$, the function $\rho^*(\boldsymbol{X}, \cdot)$ is continuous at $\boldsymbol{Z}$ for almost every $\boldsymbol{X} \in \mathcal{X}$.

2. $\rho^*(\boldsymbol{X}, \cdot)$ is dominated by an integrable function for any $\boldsymbol{X} \in \mathcal{X}$.

Then

1. $F_q^*(\boldsymbol{Z}) < \infty$ is continuous on $\mathcal{K}$.

2. $F_n^*(\boldsymbol{Z}) \to F_q^*(\boldsymbol{Z})$ almost surely and uniformly in $\boldsymbol{Z} \in \mathcal{K}$.

The assumption that the sample partitions are drawn i.i.d. is given as stated at the beginning of the proof. The first assumption of [17], Theorem 7.48 is satisfied by Lemma B.3. Since $\rho^*(\boldsymbol{X}, \cdot)$ is continuous on the compact set $\mathcal{X}$ for every $\boldsymbol{X} \in \mathcal{X}$, it attains its minimum and maximum value. Thus, there is a constant $K > 0$ such that $|\rho(\boldsymbol{X}, \cdot)| \leq K$ for all $\boldsymbol{X} \in \mathcal{X}$. This shows that $\rho^*(\boldsymbol{X}, \cdot)$ is dominated by an integrable function for any $\boldsymbol{X} \in \mathcal{X}$. Then from [17], Theorem 7.48 follows the second and third assumption of [17], Theorem 5.3.

3. Let $n \geq 1$. From continuity of F_n^* and compactness of $\mathcal{X}$ follows that F_n^* attains it minimum value. Hence, the set $\mathcal{F}_n^* \subseteq \mathcal{X}$ is non-empty. In addition, we have $\mathcal{F}_n^* \subseteq \mathcal{X}$. This shows the fourth assumption.

4. Since all assumptions of [17], Theorem 5.3 hold, both assertions hold:

1. $D^*\big(\mathcal{F}_n^*, \mathcal{F}_q^*\big) \to 0$ almost surely.
2. $V_n^* \to V_q^*$ almost surely.

5. It remains to show the Lemmata used in the proof.

Lemma B.2. *The natural projection $\pi : \mathcal{X} \to \mathcal{P}$ is continuous.*

Proof. Let $\boldsymbol{Z} \in \mathcal{X}$. From the definition of δ_2 follows

$$\delta_2\Big(\pi(\boldsymbol{X}), \pi(\boldsymbol{Z})\Big) \leq \|\boldsymbol{X} - \boldsymbol{Z}\|$$

for all $\boldsymbol{X} \in \mathcal{X}$. Thus, π is continuous. ■

Lemma B.3. *The pullback f^* of a continuous function $f : \mathcal{P} \to \mathbb{R}$ is continuous.*

Proof. Since π is continuous by Lemma B.2 and f is continuous by assumption, the pullback $f^* = f \circ \pi$ is continuous. ■

Step 3: Projecting the Solution Back

We show that

1. $D^*\big(\mathcal{F}_n^*, \mathcal{F}_q^*\big) \to 0$ a.s. $\Rightarrow$ $D(\mathcal{F}_n, \mathcal{F}_q) \to 0$ a.s.
2. $V_n^* \to V_q^*$ a.s. $\Rightarrow$ $V_n \to V_q$ a.s.

1. We have

$$\begin{aligned} D^*\left(\mathcal{F}_n^*, \mathcal{F}_q^*\right) &= \sup_{\boldsymbol{M}_n \in \mathcal{F}_n^*} \inf_{\boldsymbol{M} \in \mathcal{F}_q^*} \|\boldsymbol{M}_n - \boldsymbol{M}\|_2 \\ &\geq \sup_{\boldsymbol{M}_n \in \mathcal{F}_n^*} \inf_{\boldsymbol{M} \in \mathcal{F}_q^*} \delta_2\left(\pi\left(\boldsymbol{M}_n\right), \pi\left(\boldsymbol{M}\right)\right) \\ &= \sup_{M_n \in \mathcal{F}_n} \inf_{M \in \mathcal{F}_q} \delta_2\left(M_n, M\right) \\ &= D(\mathcal{F}_n, \mathcal{F}_q)\,. \end{aligned}$$

From $D(\mathcal{F}_n, \mathcal{F}_q) \leq D^*\left(\mathcal{F}_n^*, \mathcal{F}_q^*\right)$ follows the first implication.

2. Since F_n^* and F_q^* are pullbacks of F_n and F_Q, we have $V_n^* = V_n$ and $V_q^* = V_Q$. From this observation follows the second assertion.

B.3. Proof of Theorem 3.4

The proof consists of the same three-step procedure as the proof of Theorem 3.3. The first step can be taken from the proof of Theorem 3.3.

We show the second step: Since $\mathcal{X}$ is compact and ρ^* is continuous in both arguments, we find that ρ^* is Lipschitz on $\mathcal{X}$ and p-integrable. Since $\mathcal{S}_n$ is drawn i.i.d., we have

$$\frac{V_n^* - V^*}{\sqrt{n}} \xrightarrow{d} \inf_{\boldsymbol{M} \in \mathcal{F}_q *} \mathcal{N}(0, \sigma^2(\boldsymbol{M}))$$

by [17], Theorem 5.7. Finally, the back projection follows the same line of argumentation as in the proof of Theorem 3.3.

B.4. Proofs of Statements in Section 4

Proposition B.4. *Let $X \in \mathcal{P}$ be a partition. Then $\boldsymbol{X}^T\boldsymbol{X} = \boldsymbol{X'}^T\boldsymbol{X'}$ for all representations $\boldsymbol{X}, \boldsymbol{X'} \in X$.*

Proof. From $\boldsymbol{X}, \boldsymbol{X'} \in X$ follows that there exists a permutation matrix $\boldsymbol{P} \in \Pi$ such that $\boldsymbol{X'} = \boldsymbol{P}\boldsymbol{X}$. Then we have

$$\boldsymbol{X'}^T\boldsymbol{X'} = (\boldsymbol{P}\boldsymbol{X})^T\boldsymbol{P}\boldsymbol{X} = \boldsymbol{X}^T\boldsymbol{P}^T\boldsymbol{P}\boldsymbol{X} = \boldsymbol{X}^T\boldsymbol{X}. \qquad \blacksquare$$

Proposition B.5. *Let $\boldsymbol{C}_X = (c_{rs})$ be the compatibility matrix of partition $X \in \mathcal{P}$. The following statements hold for all data points $z_r, z_s \in \mathcal{Z}$:*

1. $0 \leq c_{rs} \leq 1$.
2. *If $X \in \mathcal{P}^+$ is a hard partition, then*

$$c_{rs} = \begin{cases} 1 & : \quad z_r \sim_X z_s \\ 0 & : \quad z_r \nsim_X z_s \end{cases}.$$

Proof. Let $\boldsymbol{X} = (x_{kj})$ be a representation of X.

1. Since all elements x_{kj} of matrix $\boldsymbol{X}$ are non-negative, we have $0 \leq c_{rs}$. From $\boldsymbol{X}^T \mathbf{1}_\ell = \mathbf{1}_m$ follows $\langle \boldsymbol{x}_{:r}, \mathbf{1}_\ell \rangle = 1$ for all columns of $\boldsymbol{X}$. Then we have

$$c_{rs} = \langle \boldsymbol{x}_{:r}, \boldsymbol{x}_{:s} \rangle \leq \|\boldsymbol{x}_{:r}\|_2 \, \|\boldsymbol{x}_{:s}\|_2 \, .$$

Observe that all elements x_{kj} of $\boldsymbol{X}$ are from the interval $[0, 1]$. Thus, we find

$$\|\boldsymbol{x}_{:r}\|_2^2 = \langle \boldsymbol{x}_{:r}, \boldsymbol{x}_{:r} \rangle \leq \langle \boldsymbol{x}_{:r}, \mathbf{1}_\ell \rangle = 1.$$

This implies $c_{r,s} \leq 1$.

2. The columns of $\boldsymbol{X}$ are standard basis vectors from $\mathbb{R}^\ell$, because X is a hard partition. From this follows the second assertion. ■

Proof of Prop. 4.1

1. Let $\boldsymbol{C}_X = (c_{rs})$ be the compatibility matrix of hard partition $X \in \mathcal{P}^+$. From Prop. B.5 follows that

$$c_{rs} = \begin{cases} 1 & : \quad z_r \sim_X z_s \\ 0 & : \quad z_r \nsim_X z_s \end{cases} .$$

for all $z_r, z_s \in \mathcal{Z}$. Suppose that $Y \in \mathcal{P}^+$ is another hard partition with compatibility matrix $\boldsymbol{C}_Y = (d_{rs})$. Then m_{11} is of the form

$$m_{11} = \chi\left(\boldsymbol{C}_X, \boldsymbol{C}_Y\right) = \sum_{r=1}^{m} \sum_{s=r+1}^{m} c_{rs} d_{rs}.$$

We have $c_{rs} d_{rs} = 1$ if and only if $z_r \sim_X z_s$ and $z_r \sim_Y z_s$. This shows that m_{11} counts all elements $\{z_r, z_s\} \in \mathcal{Z}^{[2]}$ that satisfy $z_r \sim_X z_s$ and $z_r \sim_Y z_s$. Thus, $m_{11} = m_{11}^+$. For m_{10}, we have

$$m_{10} = \chi\left(\boldsymbol{C}_X, \mathbf{1} - \boldsymbol{C}_Y\right) = \sum_{r=1}^{m} \sum_{s=r+1}^{m} c_{rs}(1 - d_{rs}).$$

We have $c_{rs}(1 - d_{rs}) = 1$ if and only if $z_r \sim_X z_s$ and $z_r \nsim_Y z_s$. This shows that m_{10} counts all elements $\{z_r, z_s\} \in \mathcal{Z}^{[2]}$ that satisfy $z_r \sim_X z_s$ and $z_r \nsim_Y z_s$. Thus, $m_{10} = m_{10}^+$. The relationships $m_{01} = m_{01}^+$ and $m_{00} = m_{00}^+$ follow similarly. This shows the first assertion.

2. Let $X \in \mathcal{P}$ be a partition. From Prop. B.5 follows that $0 \leq c_{rs} \leq 1$ for all elements of the compatibility matrix $\boldsymbol{C}_X = (c_{rs})$. Therefore all elements of $\boldsymbol{C}_X$ and $\mathbf{1} - \boldsymbol{C}_X$ are non-negative. Then the second assertion follows by definition of the function χ.

3. For symmetric $(m \times m)$-matrices $\boldsymbol{A}$ and $\boldsymbol{B}$, we can rewrite $\chi(\boldsymbol{A}, \boldsymbol{B})$ by

$$\chi(\boldsymbol{A}, \boldsymbol{B}) = \frac{1}{2} \langle \boldsymbol{A} - \boldsymbol{D}_{\boldsymbol{A}}, \boldsymbol{B} - \boldsymbol{D}_{\boldsymbol{B}} \rangle,$$

where $\boldsymbol{D}_{\boldsymbol{A}}$ and $\boldsymbol{D}_{\boldsymbol{B}}$ denote the diagonal matrices of $\boldsymbol{A}$ and $\boldsymbol{B}$, respectively. For every partition $X \in \mathcal{P}$, let $\boldsymbol{D}_X$ denote the diagonal matrix of the compatibility matrix $\boldsymbol{C}_X$. By definition,

we have

$$\begin{aligned}\mu &= m_{11} + m_{10} + m_{01} + m_{00}\\ &= \chi(\boldsymbol{C}_X, \boldsymbol{C}_Y) + \chi(\boldsymbol{C}_X, \mathbf{1} - \boldsymbol{C}_Y) + \chi(\mathbf{1} - \boldsymbol{C}_X, \boldsymbol{C}_Y) + \chi(\mathbf{1} - \boldsymbol{C}_X, \mathbf{1} - \boldsymbol{C}_Y)\\ &= \frac{1}{2}\langle \boldsymbol{C}_X - \boldsymbol{D}_X, \boldsymbol{C}_Y - \boldsymbol{D}_Y\rangle\\ &\quad + \frac{1}{2}\langle \boldsymbol{C}_X - \boldsymbol{D}_X, \mathbf{1} - \boldsymbol{C}_Y - \boldsymbol{D}_{\mathbf{1}-\boldsymbol{C}_Y}\rangle\\ &\quad + \frac{1}{2}\langle \mathbf{1} - \boldsymbol{C}_X - \boldsymbol{D}_{\mathbf{1}-\boldsymbol{C}_X}, \boldsymbol{C}_Y - \boldsymbol{D}_Y\rangle\\ &\quad + \frac{1}{2}\langle \mathbf{1} - \boldsymbol{C}_X - \boldsymbol{D}_{\mathbf{1}-\boldsymbol{C}_X}, \mathbf{1} - \boldsymbol{C}_Y - \boldsymbol{D}_{\mathbf{1}-\boldsymbol{C}_Y}\rangle.\end{aligned}$$

We set $\boldsymbol{E}_X = \boldsymbol{C}_X - \boldsymbol{D}_X$ and $\boldsymbol{E}_Y = \boldsymbol{C}_Y - \boldsymbol{D}_Y$. Observe that $\boldsymbol{D}_{\mathbf{1}-\boldsymbol{A}} = \boldsymbol{I} - \boldsymbol{D}_{\boldsymbol{A}}$ for every $(m \times m)$-matrix $\boldsymbol{A}$, where $\boldsymbol{I}$ denotes the identity matrix. Then

$$\begin{aligned}2\mu &= \langle \boldsymbol{E}_X, \boldsymbol{E}_Y\rangle + \langle \boldsymbol{E}_X, \mathbf{1} - \boldsymbol{I} - \boldsymbol{E}_Y\rangle + \langle \mathbf{1} - \boldsymbol{I} - \boldsymbol{E}_X, \boldsymbol{E}_Y\rangle + \langle \mathbf{1} - \boldsymbol{I} - \boldsymbol{E}_X, \mathbf{1} - \boldsymbol{I} - \boldsymbol{E}_Y\rangle\\ &= \langle \boldsymbol{E}_X, \boldsymbol{E}_Y\rangle + \langle \boldsymbol{E}_X, \boldsymbol{F} - \boldsymbol{E}_Y\rangle + \langle \boldsymbol{F} - \boldsymbol{E}_X, \boldsymbol{E}_Y\rangle + \langle \boldsymbol{F} - \boldsymbol{E}_X, \boldsymbol{F} - \boldsymbol{E}_Y\rangle,\end{aligned}$$

where $\boldsymbol{F} = \mathbf{1} - \boldsymbol{I}$. Expanding and summarizing gives

$$\begin{aligned}2\mu &= \langle \boldsymbol{E}_X, \boldsymbol{E}_Y\rangle + \langle \boldsymbol{E}_X, \boldsymbol{F}\rangle - \langle \boldsymbol{E}_X, \boldsymbol{E}_Y\rangle + \langle \boldsymbol{F}, \boldsymbol{E}_Y\rangle - \langle \boldsymbol{E}_X, \boldsymbol{E}_Y\rangle\\ &\quad + \langle \boldsymbol{F}, \boldsymbol{F}\rangle - \langle \boldsymbol{F}, \boldsymbol{E}_Y\rangle - \langle \boldsymbol{F}, \boldsymbol{E}_X\rangle + \langle \boldsymbol{E}_X, \boldsymbol{E}_Y\rangle\\ &= \langle \boldsymbol{F}, \boldsymbol{F}\rangle.\end{aligned}$$

Note that $\|\mathbf{1}\|_2^2 = m^2$, $\|\boldsymbol{I}\|_2^2 = m$ and $\langle \mathbf{1}, \boldsymbol{I}\rangle = \langle \boldsymbol{I}, \boldsymbol{I}\rangle = m$. Thus,

$$\begin{aligned}\langle \boldsymbol{F}, \boldsymbol{F}\rangle = \langle \mathbf{1} - \boldsymbol{I}, \mathbf{1} - \boldsymbol{I}\rangle &= \|\mathbf{1}\|_2^2 - 2\langle \mathbf{1}, \boldsymbol{I}\rangle + \|\boldsymbol{I}\|_2^2\\ &= m^2 - 2m + m = m(m-1).\end{aligned}$$

This shows $\mu = m(m-1)/2 = N$.

■

Proof of Prop. 4.2

The function χ as an inner product of strictly upper triangular matrices is continuous in both arguments. The composition of continuous functions is continuous. Critical points are points where the denominator of $\rho_1 - \rho_5$ becomes zero. These points form a set of Lebesgue measure zero. Continuity of ρ_3 follows from the assumption that $m > 1$.

Proof of Prop. 4.3

1. Suppose that $X, Y \in \mathcal{P}^+$ are hard partitions with representations $\boldsymbol{X}$ and $\boldsymbol{Y}$, respectively. The rows $\boldsymbol{x}_{p:}$ of $\boldsymbol{X}$ are binary vectors satisfying

$$x_p^+ = |\mathcal{C}_p(X)| = \sum_{j=1}^{m} x_{pj} = \langle \boldsymbol{x}_{p:}, \mathbf{1}_m\rangle = x_p.$$

Since y_q^+ and y_q only differ from x_p^+ and x_p in notation, we also have $y_q^+ = y_q$. Finally, we have

$$z_{pq}^+ = |\mathcal{C}_p(X) \cap \mathcal{C}_q(Y)| = \sum_{j=1}^{m} x_{pj} \cdot y_{qj} = \langle \boldsymbol{x}_{p:}, \boldsymbol{y}_{q:} \rangle = z_{pq}.$$

2. The second assertion follows from the fact that representation matrices have non-negative elements only.

3. We have

$$\sum_{p=1}^{\ell} x_p = \sum_{p=1}^{\ell} \langle \boldsymbol{x}_{p:}, \mathbf{1}_m \rangle = \sum_{p=1}^{\ell} \sum_{j=1}^{m} x_{pj} = m,$$

where the last equality follows from $\boldsymbol{X}^T \mathbf{1}_\ell = \mathbf{1}_m$. Obviously, the same holds for $\sum_q y_q$. ■

References

[1] R. Bhattacharya and V. Patrangenaru. Large sample theory of intrinsic and extrinsic sample means on manifolds I. *Annals of Statistics*, 1–29, 2003.

[2] A. Bhattacharya and R. Bhattacharya. *Nonparametric Inference on Manifolds with Applications to Shape Spaces*. Cambridge University Press, 2012.

[3] G. E. Bredon. Introduction to Compact Transformation Groups. Elsevier, 1972.

[4] E. Dimitriadou, A. Weingessel, and K. Hornik. A Combination Scheme for Fuzzy Clustering. *Advances in Soft Computing*, 2002.

[5] C. Domeniconi and M. Al-Razgan. Weighted cluster ensembles: Methods and analysis. *ACM Transactions on Knowledge Discovery from Data*, 2(4):1–40, 2009.

[6] V. Filkov and S. Skiena. Integrating microarray data by consensus clustering. *International Journal on Artificial Intelligence Tools*, 13(4):863-880, 2004.

[7] L. Franek and X. Jiang. Ensemble clustering by means of clustering embedding in vector spaces. *PatternRecognition*, 47(2):833–842, 2014.

[8] M. Fréchet. Les éléments aléatoires de nature quelconque dans un espace distancié. *Annales de l'institut Henri Poincaré*, 215–310, 1948.

[9] R. Ghaemi, N. Sulaiman, H. Ibrahim, and N. Mustapha. A Survey: Clustering Ensembles Techniques. *Proceedings of World Academy of Science, Engineering and Technology*, 38:644–657, 2009.

[10] A. Gionis, H. Mannila, and P. Tsaparas. Clustering aggregation. *ACM Transactions on Knowledge Discovery from Data*, 1(1):341–352, 2007.

[11] B.J. Jain. Geometry of Graph Edit Distance Spaces. *arXiv: 1505.08071*, 2015.

[12] B.J. Jain. Properties of the Sample Mean in Graph Spaces and the Majorize-Minimize-Mean Algorithm. *arXiv:1511.00871*, 2015

[13] U. Lang *Length Spaces*. Lecture Notes, ETH Zürich, 2004.

[14] T. Li, C. Ding and M.I. Jordan. Solving consensus and semi-supervised clustering problems using nonnegative matrix factorization. *IEEE International Conference on Data Mining*, 2007.

[15] M. Meilă. Comparing Clusterings – an Information based Distance. *Journal of Multivariate Analysis*, 98(5):873–895, 2007.

[16] J.G. Ratcliffe. *Foundations of Hyperbolic Manifolds*. Springer, 2006.

[17] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming*. Society for Industrial and Applied Mathematics, 2009.

[18] A. Strehl and J. Ghosh. Cluster Ensembles – A Knowledge Reuse Framework for Combining Multiple Partitions. *Journal of Machine Learning Research*, 3:583–617, 2002.

[19] A.P. Topchy, M.H. Law, A.K. Jain, and A. Fred. Analysis of consensus partition in cluster ensemble. *IEEE International Conference on Data Mining*, 2004.

[20] A.P. Topchy, A.K. Jain, and W. Punch. Clustering ensembles: Models of consensus and weak partitions. *IEEE Transactions in Pattern Analysis and Machine Intelligence*, 27(12):1866–1881, 2005.

[21] S. Vega-Pons, J. Correa-Morris and J. Ruiz-Shulcloper. Weighted partition consensus via kernels. *Pattern Recognition*, 43(8):2712–2724, 2010.

[22] S. Vega-Pons and J. Ruiz-Shulcloper. A survey of clustering ensemble algorithms. *International Journal of Pattern Recognition and Artificial Intelligence*, 25(03), 337-372, 2011.

[23] N.X. Vinh, J. Epps, and J. Bailey. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *Journal of Machine Learning Research*, 11:2837–2854, 2010.